Introduction

Isn't it surprising that exceedingly bright physicians, technologists, and salespeople who are involved daily with computing devices such as CT & MR scanners and digital subtraction angiography equipment, can be so deficient of knowledge and even timid when it comes to every day desktop computers? Over the past few years, it has almost become impossible to pick up a newspaper or magazine without reading a headline about desktop computing. And when it comes to radiology, I have heard it said that at the RSNA you can't even sell barium bags anymore if they're not hooked up to a workstation! Yet many radiologists, who make their living at the console of a million-dollar CT, MR or angiography display terminal, are completely unfamiliar with or terrified by a PC.

In the next two years we will witness the explosion of digital imaging devices and networks in the hospital, home and even the grocery store. If you can run a CT console or angio table but don't feel comfortable with a PC, workstation, or archive this book is for you. This is no better a time to refrain from learning as from the challenge of learning cross-sectional imaging during the 1980s. If you are timid, uncertain, and insecure, read on... In these pages you will find a simple explanation for everything ranging from a mouse to a digital tape jukebox. When you finish you should be able to trade bits and terabytes with any pencil neck computer nerd. (If you can't, e-mail me at berman@lumisys.com).

Phillip Berman, M.D.
President, Lumisys, Inc.
April 18th 1998

P.S. It is with great pride that I can report that over the past 12 months more than 1,000 copies of this Primer have been requested and distributed by and to radiologists, technologists, industry product managers, consultants, and the like. We'd be delighted to make additional copies of the PACS Primer available to your organization and/or customers. If you're interested, please call Vicki Peterson at (520) 298-1000 ext. 111 for details.

The Workstation

The "workstation" is basically a PC on steroids. Usually they cost more, specifically because they contain a great deal more RAM (expensive) and support more than one video card/monitor (also expensive).

The workstation or PC comes with some basic components. First there's the CPU (Central Processing Unit). The CPU is the hardware that greatly determines pricing. For instance, Intel makes 80386 (aka 386), 80486 (aka 486) and 80586 (aka Pentium), now there are 80686 (by Cyrix), Pentium Pro (related to the 586) and Pentium II (a Pentium Pro with MMX instruction set). Motorola makes a whole bunch of different CPUs, which power the Macintosh series (I hesitate to name them here because they change every month or so). Sun makes CPUs for Sun Workstations, Silicon Graphics for SGI workstations. The price of the CPU (and the other goodies "on board") essentially determines the speed and affordability of the computer.

Then, there's what's called an operating system (or OS). Examples are DOS, UNIX, Windows'95, Windows NT, Macintosh system 7.5, and VMS. Your computer is a lump of silicon and metal without an OS. The operating system is what made Bill Gates important and very wealthy. The operating system is what determines what software you can and can't run on the computer. If you have UNIX, forget running Microsoft Word for Windows. If you have Windows'95, you can't run Macintosh software.

Operating systems are usually designed to run on a specific CPU or series of CPUs. That is, you can run DOS on an Intel 286, 386, 486 or Pentium, but forget about running it on your IBM AS/400.

The operating system also is responsible for the general look and feel of the computer when you sit at the screen. For instance, DOS looks like a blank black screen with a bunch of incomprehensible characters like:

C:\>

And you have to type incomprehensible commands like "cd\bin" to talk to the computer. As a result, many manufacturers hard coded buttons on their CT and other imaging computers with options like "Next Image" and "Rotate Image" and "ROI" so you didn't have to know how to speak geekish in order to do your job.

As computers have become more powerful and compact, operating systems have been designed to take advantage of this power and become graphical so that ordinary people can interact with the computer more intuitively. These are called GUIs (or goo-eys) for Graphical User Interfaces. The Macintosh was probably the first widely accepted GUI. Then came the Microsoft GUI (aka Windows) -- again making Bill Gates important. (Starting to get the trend here?)

Once the first GUI took hold, more GUIs followed in their footsteps. While the workstation crowd doesn't like to admit it, now you can even find major medical equipment, including MR and CT scanners and virtually every kind of PACS component running GUIs on PCs.

The "next generation" of GUIs is the OOUI (Object Oriented User Interface, pronounced, "OO-EE"). OS/2 is an OOUI; Windows 95 is an OOU; so are NextStep and the Mac. While the distinction between GUIs and OOUIs may not be clear at first, it becomes more intuitively obvious if you run the desktop. The GUI is application centered; the icons are used to open applications but cannot otherwise be manipulated. After you start an application, you can select objects (like documents or images) to run within it. Menus are the method of navigation, and objects are created in an applications-specific manner.

OOUIs, on the other hand, have icons that can be manipulated without being opened. For instance, you can take a document's icon and drop it in the waste basket, on a printer, or on another networked desktop and it will do what you would expect -- you don't need to run the word processing application to do something to it. A window provides a view of what's inside an object -- there is a one-to-one relationship between the window and the object. A template folder allows you, and in fact encourages you, to create new instances of an object and/or new objects. OOUIs provide container-type objects (folders and notebooks). These in turn can act on other objects; actions performed on container objects affect all the objects inside the container. Theoretically all the applications behave in the same way, and the user is the visual programmer of the desktop.

All this may seem confusing, but if you re-read it after playing on an OS/2, Windows 95, or NT 4.0 system, you will understand. The most peculiar thing about OOUIs is that experienced GUI users have difficulties when they switch to them, whereas new users find them much easier to learn by all accounts. It is believed that because GUIs are a poor abstraction of the real world, it takes some re-orientation for GUI users to get used to something more closely resembling the paper world paradigm.

An excellent example of this is the OOUI abstraction in our DICOM 3.0 workstation, *i*VIEW Pro. This program contains no menus, no keyboard commands, or a disappearing toolbar - in short, nothing but image *objects*. Since films and images are like objects in the real world, radiologists actually 'get it' pretty quickly (lucky for us).

A. Hard Drives

There's also other stuff in that little box. The hard drive (Winchester Drive if you're easily impressed by long, proper names) is the thing where images, documents, the operating system, and all the important stuff are contained. These things are stored as files just as physical files are stored in a filing cabinet. Rest assured it'll still be there after you turn your computer off, then turn it on again: it's almost like leaving the office, turning out the lights and coming back the next day. Except, hard drives crash and burn! One day you'll turn the machine on again and get a message like "Your hard drive can't be found," or "No hard drive found."

Everyone is always warned: "Never forget this: back up your data." Having read that, you won't take it seriously until your hard drive melts down for the first time. The miserable thing is that hard drives only crash every year or two on average, so you will almost certainly be lulled into a false sense of security before it happens to you.

This is why only the people who are on their second computer have backup tape drives, backup hard drives, or system administrators. No one ever learns this the easy way, and don't even begin to think you're an exception.

Hard drives come in various sizes, and the bigger they are the harder they crash. They are measured in megabytes, or MB for short. In 1985, 20 or 40 MB was a pretty good size disk, but today, the Windows operating system alone sucks up 30 MB. In 1993, 200-300 MB was pretty average and cost about $300-400. In 1995, 1 gigabyte (1 GB = 1,000 MB) was equally ordinary and cost about $400-500. Today, 6 GB hard drives are not unusual, and the cost remains in the same neighborhood. Get the trend? (Hint: more hard drive space costs about the same over time.)

Saving stuff is what computers are all about. But hard drives are physical things, so putting a second one in your computer isn't something most of us like to do. In order to conserve space (and reduce transmission times for images sent over phone lines), PACsters like to "compress" the images into smaller files. Also, PACsters like to delete files after they aren't used for a while. This is fine, except that the disk will eventually become "fragmented" because of the way these little Winchesters work. You'll need a "de-fragmenter" program to fix this problem. Like most things, using a defragmenter takes time and is extremely boring, so you'll either need a de-fragmenter that works while you're not there (at night) or a system administrator (a human).

Some hard drives are faster than others. The time it takes for your drive to read an image, decompress it if it's compressed, then put the image on your screen, depends on the drive's speed. "Disk access time," measured in milliseconds, is important to your workstation's performance, particularly with large image files. Various types of disk controllers exist, and they play a large role in determining how fast the access time is. IDE (Integrated Device Electronics), EIDE (Extended IDE), and SCSI (Small Computer Serial Interface, pronounced "scuzzy"), and superwide SCSI 2 are most popular. For the most part the SCSIs are the fastest and most highly preferred.

Additionally, to conserve space and reduce transmission time when using telecommunications, images on hard disks are often "compressed" by mathematical algorithms (JPEGs, wavelets and other animals, which we will discuss a little later). Suffice it to say for now that reading the images from the hard disk requires that they be "uncompressed" before they can be displayed on your video monitor. No matter which technique is used, decompression takes time... too much time for you and me.

B. RAM

So, in general, hard drives, no matter how fast, are too slow for all but the most brain-dead radiologist's needs. That's why they made random access memory, or RAM. The "random access" part of RAM means that (unlike on your hard disk) any piece of information can be located as swiftly as any other, no matter where it resides.

RAM is also volatile; that is whatever it contains evaporates into the ether when you turn off your machine. RAM is where your images and documents live while you are working on them. If you haven't saved them to the hard disk and your local Electric company has a bad couple of milliseconds of service, your work will vanish in that millisecond. Poof! But, fortunately for you and me, unlike documents, this tends not to be a problem with PACS because images are usually first saved to the hard disk and then displayed by "loading" them into RAM.

RAM is fast. Like most fast things it is expensive. It is measured in megabytes, just like hard disks. A computer from the Price Club will likely have at least 16 megabytes of RAM, but if you use your PC at home for teleradiology, you probably need 32 or 64 megabytes. If you buy a workstation from a major medical manufacturer it will likely have 64 or 128 megabytes of RAM. At 1997 prices, RAM is worth $3-$5 per megabyte. Upgrading your PC to a so-called workstation (from 16 to 64 MB) will run about $200 in hard costs; then you may need to add installation costs and re-seller margins.

So by the time everyone finishes picking through your wallet, unless you buy it wholesale and open the box yourself, such an upgrade will cost you more than the computer probably cost from Dell, Gateway, Price Club or Joe's computer store in the first place. Like most fast and expensive things, we all want more than we can usually afford.

C. Rodents & Other Pointing Devices

Then there's the mouse... the funny little thing with a cord, and one, two or three buttons on it. (One for the Macintosh, two for Bill Gates, three for people who are adroit enough to use their middle fingers while computing). This is what allows you to point and click at stuff instead of always using a keyboard or keypad. With it you spend more time looking at the screen and less looking at your fingers. They also make trackballs, which are basically, mice on their backs. If you turn a mouse over you'll see a small little ball which needs to be cleaned every so often (especially if your kids leave gum around your computer). That ball keeps track of where the cursor (arrow) resides in the 2D space of the image monitor. A trackball simply puts the ball on top and allows your fingers to move it directly, keeping the device from wandering around the desktop, instead of pushing a little plastic mouse around. The buttons are arranged around the ball.

There are also touch screens where the position sensor is laid on top of the video screen. This is way cool at first but a nuisance after a short while. One reason for this is that the screen gets filthy. Another is that you get (or give) other people germs from unwashed hands in a medical environment (remember the germ theory?).

Then there are joysticks and light pens, which aren't exactly ubiquitous. Just like everywhere else, when 5 or 6 options are present in a marketplace, the best solution has yet to be offered. It isn't too hard to figure out what the best solution would be (hint: it's your voice); unfortunately, it's much harder to offer it. As they say, maybe at next year's RSNA...

D. Mother, Please Take the Bus

Inside your computer is a "motherboard." This board basically is the computer, to which is attached a power supply, surrounding metal case and miscellaneous devices (like the keyboard, mouse and monitor). On the motherboard sits the CPU, or chip (for example your Pentium II chip running at 300 megahertz), the RAM, and connections to the hard disk. The motherboard has connectors or slots to attach other goodies to your computer. If you ever open a computer, you'll see a bunch of empty slots waiting for you to fill them. The number of slots varies among computer motherboards and has a large impact on the expandability of a computer. (Without opening the case you can sort of guess how many of these slots there are by looking at the back of the computer. You will see a number of little metal slider sheets blocking your view of the computer's innards. Each one of those is an empty slot.)

You will fill one slot with a video display board. This is how your computer tells the monitor what to display. Basically, if the information is an image, it is taken from the hard disk (and probably decompressed), written into the system's RAM, and then taken from system RAM and written to the video display card. The display card then displays the information to the monitor through a cable connected to it from the back of your computer. (To make things more complicated, the video display card also has RAM that is more expensive and less expandable than the “system” RAM already discussed above. For the most part, 8 or 16 MB of video card RAM is about all you can buy on a commercial card).

All this traffic must go back and forth quickly if your image is going to make it to the screen in the blink of an eye. Each time the information moves, it takes time. Additionally, the "pipe" through which it moves can be narrow or wide and is known as "the bus". There are lots of buses, but for the most part the CPU and operating system combination function on a predetermined or specified bus. In other words, once you buy the computer, you can't change the bus.

Because the bus is a bottleneck, a new bus is designed from time to time to speed performance. Since the original IBM PC, the ISA bus has been reworked a few times for just that purpose. First an extended ISA bus was designed (EISA), but it wasn't terribly popular because it increased the cost of the motherboard and improved performance to only a small share of the desktop market. Then the VESA local video bus was designed to improve video performance - an improvement needed by many who were running video intensive Windows applications. Recently the PCI bus has made its appearance for similar purposes and seems to have made it as a standard.

Additionally, a shortcut can be taken by PACsters to achieve better performance: refuse to take the bus. With special engineering, jumper cables can be applied so that the data can bypass the bus on the motherboard and go directly to the video card from the hard disk. Another trick is to isolate the image traffic on a "backplane" (i.e. a second and separate) bus from the rest of the general computing traffic. A third alternative is to use accelerator cards.

Any of these choices, properly engineered, will improve performance, but special designs and additional hardware like this create their own set of problems: First, the computer architecture is immediately and inherently proprietary. Second, price quickly increases with the addition of special hardware, additional boards, etc. Third, maintenance can become difficult for anyone but a system engineer from the manufacturing company. Upgradeability can become questionable after a few years. Adding a second, third, fourth or eighth monitor may not be possible or may require additional accelerator boards or a second, third or fourth image bus (i.e. the solution may not "scale" well). Interconnectivity between various systems may not be possible.

E. Bits Are Not Bytes

Technical stuff: A bit is the simplest form of data. It's either a 0 or a 1. A byte is 8 of these bits strung together into a "word" something like 10101110. (It's one example of a word, but some computer words are 16 bits or 32 bits long.)

F. My Modem's Faster than Your Modem...

Modems are the screechy little beasts that send information over phone lines. They are what make that intolerable sound between FAX machines. They dial, connect, handshake and then start transmitting. Most of the screeching you hear is the handshake, and it's about as weird and unique as those high-fives which follow long touchdown runs in the NFL. Modem stands for MOdulation-DEModulation. And that's exactly what's going on between two modems when they're connected.

A few of us can remember when modems had suction cups in which you put the telephone receiver (which had round earpieces). Those modems functioned at 150 or 300 baud, or bits-per-second. Today your modem connects directly to a modular phone jack. It may be of the "internal" or "external" flavor. An internal modem fits in one of those slots on your computer motherboard. An external modem hooks up to the serial port of your computer (which itself is a half-slot card inside your computer).

The external modem flavor has its own separate power supply and a bunch of meaningful lights which flash on and off while in use. The internal flavor sucks electricity directly from your computer, so it's usually a bit cheaper and may be faster if your serial port card isn't one of the fastest. But when you need to "reset" or "reboot" your modem (translation: computer talk for "turn off and then turn on your computer so that the problem you're experiencing goes away"), you will have to shut your whole computer off if you have the internal variety. Not so with the external variety.

Today's modems are quite a bit faster than the early ones. Current modems function at 14,400 bps to 28,800 bps. This sounds pretty fast until you do a little math. Since a baud is a BIT per second, and there are 8 bits in a byte, you only get about 3,200 bytes per second through a 28,800 bps connection. Since an uncompressed frame grabbed CT slice is about 325,000 bytes, that's about a 100-second transmission at 28,800 bps - nothing to brag about.

Things can get even slower in the real world. First, most modems are of the 14,400 bps or 28,800 bps variety. Second, even with a 28,800 bps modem, most ordinary analog phone lines poop out at somewhere between 14,400 and 24,400 bps. So that translates to 2.5 - 3.5 minutes per CT image (not exactly worth bragging rights). And that's just a frame-grabbed image, which is 8-bits, not the holy grail DICOM 16-bit image. The original data or DICOM image is 16 bits per pixel -- twice as big (~650,000 bytes per slice). That's why PACsters compress images and why there's a digital phone line somewhere in your future (aka ISDN, see below) for plain films and DICOM images if you want to be a couch potato radiologist at home at night keeping up with the Jones'.

G. Monitors

Monitors come in two shapes, just like photographs: portrait and landscape. Of course, radiographs are usually displayed in portrait format (long side displayed vertically). And, of course, computers are usually bundled with landscape monitors. So, as you might expect, the portrait monitors cost more than the landscape monitors - another tough break for the purchasing radiologist or healthcare enterprise. One solution is to buy a monitor that pivots between landscape and portrait modes. This costs even more. Another solution is to suffer with the landscape mode -- a much more popular decision for radiologists at home.

Then of course there's the monochrome/color problem. Everyone likes color except, of course, you guessed it: radiologists. (But then, people in ultrasound and nuclear cardiology like color too.) As a result, most radiologists prefer a gray-scale monitor to the ordinary color monitor. This is a good choice from the perspective of resolution, because color monitors function on the RGB (or Red-Green-Blue) principle. That is, there are three dots assigned to any one pixel: one red, one green and one blue. These dots take up physical space on the monitor and reduce the displayed pixel density of a gray-scale image. The brightness of a color monitor is poorer than that of a gray-scale monitor for the same reason.

So, you think you want a gray-scale portrait monitor? Well, think about this: Monitors are fabricated in a way that specifies an inherent maximum resolution. That is, some monitors are only capable of 480x640, some up to 600x800, some up to 748x1024, 1280x1024, 1600x 1280, 2048x1536, and finally (what you and the ACR have been waiting for) 2560x2048. The bigger the matrix, the bigger the monitor. Most 2Kx2K monitors (and there aren't that many) weigh in at about 150 pounds and are about as big as a small icebox! The bigger the display, the bigger the price tag. Monitors capable of 2Kx2.5K display usually cost in the range of $15,000 each. New flat panel technologies will probably make these statements untrue by 1998 or 1999; but for now, it’s still the refrigerator.

Then there's the problem of the video display card. Video cards are either built for landscape or portrait mode. The card must also be matched to the monitor's display frequency (i.e. they need to be synchronized), and the monitor must be able to display the matrix size of the card. A 2Kx2K card is pretty useless with a 1Kx1K monitor. And a 2Kx2K monitor is a waste of money for a 1Kx1K video card.

You also need enough RAM on the card to fit the decompressed image or images on board. So, for instance, if you have a 2Kx2K 16bit grayscale chest image, that'll be 8 megabytes of video information on an uncompressed basis. RAM costs dough. If you want more than a single monitor display, most manufacturers will need the same number of cards as monitors. So to display 2Kx2K on a 2Kx2K monitor, you will need lots of money: $15,000 for each monitor, $7-10,000 for a card for each monitor equipped with enough RAM, and you wouldn't want to skimp on the computer would you? Multiply this by as many monitor/cards as your wallet desires and you start to get chest pain.

Finally, there's the brightness/luminance issue. View boxes provide 200 foot-lamberts of luminance (don't ask me what a foot-lambert is). Most monitors you buy with your Price Club special display at 30 foot-lamberts. Most gray-scale monitors sold into the teleradiology/PACS market are presently in the 50 foot-lambert class and seem to function reasonably well in the environment. Some newer models display at 80-100 foot-lamberts. Guess what?: they cost a lot more money.

H. Image Files

Images are kept in files on a computer much like documents are. There are standard file formats that are read by the software and then displayed as described above. In the good old days, manufacturers created their own (proprietary) file formats. These were good because no one could read their image files except a piece of equipment they sold to you. The files were also compact and very fast in terms of communication over networks (in order to archive to tape, print to film, or review on a remote display station). Unfortunately, this created the radiological equivalent of Bosnia. Everyone was armed and dangerous and no one spoke the same language.

Meanwhile, the rest of the real computing world created many standard image file formats, which could be used in document and graphics publications. Many of us pencil neck geeks yearned for our medical images to be able to participate in our word processor and presentation graphics programs. Some of us even created teleradiology applications that could do just that by using cute little formats like TIFF, PCX, BMP, and GIF. Aside from the pixel information, the "header" in each of these file formats varied in format and size. To make life interesting, the people in charge of NEMA and the ACR created their own image file format, suprisingly known as the ACR/NEMA format. Now it's been renamed in its third iteration as DICOM (short for Digital Imaging and Communications in Medicine) version 3.0. In DICOM, gray scale images are 16 bits per pixel (or 2 bytes per pixel), and true color images are 24 bits per pixel plus 8 bits per pixel of intensity information (or a whopping 4 bytes per pixel).

The spatial resolution, or size, of a digital image is defined as a matrix with a certain number of pixels (or information dots) across the width of the image and down the length of the image. The more pixels, the better the resolution. This matrix also has "depth." Depth is usually measured in bits and is more commonly known as shades of gray: 6 bit images have 4 shades of gray, 7 bit images have 128 shades of gray, 8 bit images have 256 shades of gray, and 12 bit images have 4096 shades of gray.

The file size of a particular image is determined by multiplying the number of horizontal pixels "by" (or "times") the number of vertical pixels, and then by multiplying the number of bits in the shades of gray as the depth. For example, an image might have a resolution of 640x480 and 256 shades of gray, or 8 bits of gray scale depth. The number of bits in the data set can be calculated by multiplying 640 x 480 x 8 = 2,457,600 bits. Since there are 8 bits in a byte, the 640 x 480 image with 256 shades of gray is 307,200 bytes or 38,400 bytes of information.

Now, bytes can't be fractionalized. So if you have a 12 bit image and there are 8 bits in a byte, you need 2 bytes to express all the info. The last 4 bits (referred to as the "high bits") are zeroed. Still, the size of the file needs to be calculated with two bytes per pixel in the setting of 12 and 10 bit images. (N.B.: Computed Radiography uses 10 bit images for the most part.)

OK, now for a little more math: the above means, an ACR certifiable, diagnostic plain film image is a minimum of 2Kx2K , that's 2,000 by 2,000 pixels for a total of 4 million pixels. Each pixel has two bytes (or 16 bits) of info, for a total of 8 million bytes (translation 8 megabytes) per image. (The header adds a few thousand bytes, but who cares with an 8 megabyte image?). 8 megabytes buys you a lot of RAM requirements, a 150 pound monitor and hours of transmission time over ordinary phone lines. Enter compression…

I. Squashed!

Even a single 8 megabyte chest image or KUB transmitted over a 128 kilobit per second ISDN (digital) line will take 500 seconds (or 8 minutes) to transmit. A local area network running plain vanilla Ethernet (see later) functions at 10 megabits per second. At 100% efficiency (and efficiency is often no better than 35%) this would tie up the network for 6 or 7 seconds per image - a situation that most information systems people will not tolerate in a hospital or clinic.

Since the math isn't acceptable to anyone, the alternative of compressing images seems worthy. Now, one can compress images in a lossless (non-destructive) or "lossy" (destructive) way. Obviously the former is preferable, but unfortunately the maximum that most mortals can non-destructively compress an image is about 2.5:1 or 3.0:1. While that helps, it isn't enough: recalculating the above math you get 2.5 minutes per 2Kx2K image over a 128 kbps ISDN line and 2 seconds on a LAN.

Enter destructive (aka lossy) compression. Higher compression is destructive by its nature but not necessarily detectable to the eye. Compression algorithms may be symmetrical or asymmetrical. Asymmetrical algorithms take longer to compress than decompress, or vice verse (get it?). In general, algorithms that take longer to decompress are greatly disfavored because the radiologist will throw a hissy fit waiting.

Compression algorithms come in a variety of flavors. The most popular are LCZ and JPEG. The hippest and newest compression algorithms are based on wavelets. Since DICOM 3.0 only permits JPEG. JPEG is pretty good, is reasonably fast on both the compression and decompression times, and is widely implemented. Some enhanced versions of JPEG permit visually acceptable compression at ratios of 40:1 - 60:1. Now the math improves some…

Certain images can withstand a variable amount of compression without a visual difference -- CTs and MRs have large areas of black surrounding the actual patient image information in virtually every slice. The loss of some of those pixels does not affect the perceived quality of the image, nor does it significantly change the reader's interpretive performance.

Now, for realism: For 'conventional' JPEG 10:1 is reasonable for plain films, with higher ratios for CT, MR, nuclear medicine, and ultrasound. So you can get a 2Kx2K image over ISDN in 50 seconds, over 10MB Ethernet in 6 /10ths of a second (assuming the unrealistic throughputs) or 1 second (assuming realistic throughput). Enhanced JPEG (eJPEG) will support 30-70:1 compression without visible diagnostic quality degradation, improving this math by a factor of 3-7 again. This technique however works far better on plain films than on small file format images like CT, Ultrasound, MR and Nuclear Medicine (for some truly boring and technical reasons I refuse to disclose here). So on the small file formats: we're really in the 10-20:1 zone for visually acceptable compression ratios; clearly higher on plain films with eJPEG (fortunately).

The core technique of JPEG compression uses blocks of pixels (usually 8x8 pixels) and performs mathematical compression on those blocks. This is easy to detect by repetitively magnifying a compressed JPEG image on workstation (try it - its fun!). Enhanced JPEG is also called "full-frame JPEG" because the "block" is the full frame or entire image. (This is the same 'first step' used in wavelet compression techniques.) So repetitive magnification of enhanced or full-frame JPEG merely makes the image look degraded, not 'blocky'.

Then there's the murky land of wavelet compression. Wavelet mathematical techniques for compression are less than two decades old. Wavelets use frequency information for compression. Techniques vary greatly and are not yet very mature. The best techniques are probably yet to be discovered, unlike very well understood and long-standing science of JPEG compression techniques.

Many small companies have been quick to commercialize wavelet compression for use in commercial home markets. Unfortunately, a few small companies have adopted wavelet compression in teleradiology and 'spun' this somehow as part of the DICOM standard. It is clear that wavelet compression falls outside the agreed upon DICOM 3.0 standard. The net effect is that a reputable manufacturer who offers a genuinely and fully compliant DICOM 3.0 workstation may well not be able to display a wavelet compressed image. As well, the acceptable compression rates of both full-frame JPEG and wavelet techniques as implemented to date are pretty much the same as full frame (enhanced) JPEG compression rates are better for plain films and tolerable at 30-70:1 on 2kx2k images (depending on their character of course); and they are inferior but similar to JPEG compression ratios for small file format images.
The net result is no surprise: no free lunch! Nevertheless, things are moving in the right direction. And if things play out as they usually do, by the time everyone has wavelet figured out and integrated in the DICOM standard, we'll all have T1 bandwidth to our toasters and won't need any compression for our images. Now we're smokin' on the infobahn, eh?

J. Eat at Terabytes Today

One last subject for the workstation and its general connectivity is megabytes and terabytes. Hard disks are considerably cheaper and bigger these days. You can buy 6 gigabyte (that's 6,000 megabytes) hard disks for a PC today for about $400, or $0.07 per megabyte. Now at 10:1 compression that's about $0.07 per 2Kx2.5K DICOM image; at 30:1 compression that's $0.02 per image.

Enter the cheaper, networked archiving media: digital tape and optical drives. Now I don't want to get into the detailed economics of these devices because they change quickly and the pricing is never firm when you buy a $1 million digital tape jukebox, but suffice it to say that certain things are substantially better in this environment, one of them being price. Another consideration is that tape allows for redundancy, spreading more than one copy on a different physical tape, or re-writing on an automatic basis so that Mrs. Smith's studies are physically located in close proximity on the same tape. Tape tanks are also manageable in terabyte type sizes. Nondestructively compressed at 3:1, a terabyte will hold 375,000 2Kx2K DICOM images. So tape, as they say, "scales" properly.

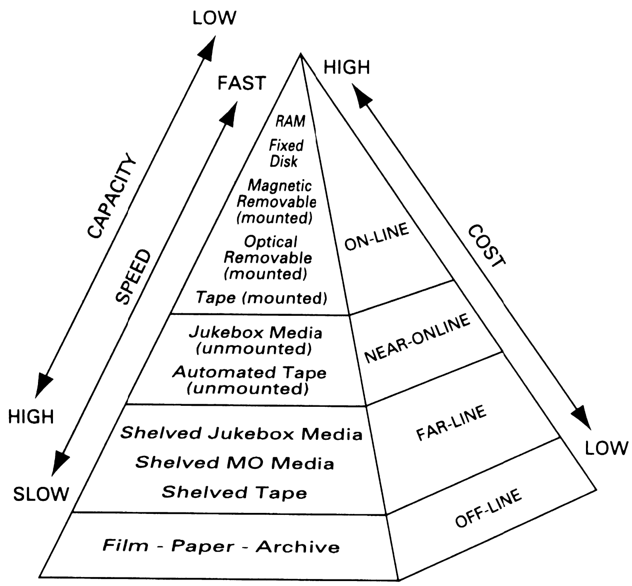

Networking 101

Telecommunications and network communications are the backbone of PACS and Teleradiology. Without them, nothing gets in or out of the workstation. We talked a bit about interaction between file size and throughput above, but now it's time to suffer a little more.

First you need to get the concept of bandwidth. This is a very cool word and once you understand it, it sets you apart from the rest of the crowd. Bandwidth is the amount of data that can be transmitted over a channel, measured in bits per second. For example, Ethernet has a 10Mbps (that's geek for Megabits per second) bandwidth and FDDI has a 155+Mbps bandwidth. (Warning: Actual throughput may be different than the theoretical bandwidth given in these numbers and *usually is by a factor of 40-60%, unfortunately on the downside*).

A. LANs

A LAN (Local Area Network) is a group of computers; each equipped with an appropriate network adapter card and software, that share applications, data, and peripherals. While all of the connections are made via cable or wireless media, a LAN does not use telephone services. If your computers are "hard wired" together, this is usually faster than a phone line or other communications line. A LAN typically spans a single building or campus and is usually wired

through a "hub" so that any workstation can connect to any other workstation or device located on the network.
There are various ways of wiring LANs. The most common is a "star" topology or shape, where each workstation has a run of wire back to "home," or the hub, known as a "home run." In a star topology network, the nodes are connected in a hub and spoke configuration to a central device or location. The "hub" is a central point of failure.

Other methods include Token Ring, which is a 4Mbps or 16Mbps network that uses a logical ring topology, but a physical star topology. A token is passed as the method by which access to broadcast on the network is allowed. Each ring can have up to 256 stations.

In addition to the "topology," there are "media," "protocols," and "network operating systems." The media are the wires. Wiring is usually in the form of copper wires, twisted and unshielded electrically, hence the name UTP (Unshielded Twisted Pair) wiring. Different classes of wiring based on gauge and electrical properties are available. They are creatively named UTP classes 1, 2, 3, 4, and 5. As a general rule, the thicker the wire, the better the quality and the higher the price.

Category 1 wiring is old-style unshielded twisted-pair telephone cable, and it is not suitable for data transmission. Category 2 UTP is suitable for snail mail rates of up to 4Mbps. Category 3 UTP can handle speeds of up to 10Mbps, and it is the minimum cable required for the ever-present, vanilla Ethernet network running "10BaseT." (10BaseT is the implementation of Ethernet on unshielded twisted-pair wiring. It uses a star topology, with stations directly connected to a multi-port hub. It runs at 10Mbps, and it has a maximum segment length of 100 meters.)

Just for fun, the wire pairs should have at least three twists per foot, but no two pairs can have the same twist pattern. Category 4 is the lowest-grade UTP acceptable for 16Mbps Token Ring. And category 5 is certified for speeds of up to 100Mbps, but 155Mbps will likely be possible when ATM is fully implemented in a LAN environment.

Category 5, as a result, is the fastest and most highly preferred at newer installations. As usual, it's also the most expensive (about $0.11/foot - this doesn't sound like much, but remember that each workstation has a home run). If you don't have at least "Cat 5" wiring, you are definitely computing in the stone ages by now.

Other kinds of media include coaxial cable (just like your TV) and fiber-optic cable. Now this stuff is real expensive (like $1/foot), but if you want PACS style throughput, fiber will likely be in your PACS' future.

Fiber Distributed Data Interface (FDDI) is the specification for a 100Mbps network that is implemented as dual, counter-rotating rings. A fiber FDDI network can support up to 500 stations over about 1 mile of fiber. FDDI-I is specified for data networking only. FDDI-II, which has yet to be finalized, will specify the transmission of both voice and data. FDDI, originally specified to run only over fiber, can also operate over shielded and unshielded twisted-pair, although the distances tolerated are greatly shortened.

Now where does all this wiring go? Usually the wire is terminated into a "jack" that looks like a modular phone plug (jack for those of you who are picky). In fact, it is a phone plug, but it is larger than the one in your house. The one in your house is likely an RJ-11 jack with four wires, and the network plug is likely an RJ-45 jack with eight wires which can be used with either 10Base T-Ethernet or a PBX. If you use it for a phone, likely the cord will look flat. If it is to be used for networking, it will likely be round with twists in it. (If you use flat cord for networking, you will cut down the throughput of your network substantially. So don't do that.)

Then there's the end you plug the wire into that resides in your computer. Again, this eats one of those slots or half-slots in your computer and is called a network interface card (NIC). Lots of people make these. They come in lots of flavors and cost about $50-150 dollars for Ethernet types. Their manufacturers are usually on the stock exchange and sell for 200 or 300 times next year's earnings. It’s safer to buy their products, in general, than shares of their stock.

You should always remember that the LAN is fast compared to most other forms of telecommunications (which are referred to as WANs, or Wide Area Networks). An easy point of reference is that a vanilla LAN running Ethernet should have a signaling rate of 10 MB per second. Even with 35% efficiency, that's 3.5MB per second. A T1 line has a signaling rate of 1.544 MB per second - only 1/3rd that speed.

B. The NOS (Network Operating System)

A network operating system (NOS) is the software that runs on a file server which governs access to the files and resources of the network by multiple users. Most hospital-size networks will have a NOS, whereas smaller offices (particularly since the advent of Windows for Workgroups, Windows'95 and Windows NT) may not.

Examples of NOSs include Banyan's VINES, Novell's NetWare, and IBM's LAN Server. Perhaps the best known and dominant market NOS is NetWare, Novell's network operating system. NetWare uses a variety of network communication protocols with enchanting names like IPX/SPX, NetBIOS, or TCP/IP network protocols. It supports DOS, Windows, OS/2, Macintosh, and Unix clients. NetWare versions 4.x and 3.x are 32-bit operating systems; NetWare 2.2 is a 16-bit operating system.

C. WANs

1. So Pick up the Damn Phone!

Plain Old Telephone Systems (or POTS) are cheap, ubiquitous analog circuits that allow throughput of up to 33,600 BITS (not bytes) per second. You will see higher numbers, but they usually include a multiplication for presumed compression applied to the data by the modem itself, or a special circumstance where modems on either end are identical and the POTS line supports the higher rate (not very common). For the most part, except in unusual circumstances, POTS don't really tolerate much more than 19,200 - 22,600 throughput. Such lines are incapable of real time 30 frame/second video (even at very poor resolution) and are terribly slow for most PACS and uncompressed Teleradiology applications. They have however been quite effective for compressed, on-call teleradiology applications, primarily for digital modalities (e.g. CT, ultrasound, MR,

nuclear medicine) and for plain films when compressed by enhanced JPEG and (non DICOM) wavelet methods. [More about this later.]
You will no doubt hear a bunch of letters and numbers rattled off by those in the know. They will say stuff like "v-dot-32-bis". Here's what that all means: A V.32 bis modem offers speeds in increments of 4800bps, 7200bps, 9600bps, 12,000bps, and 14,400bps. A v.34 bis modem will go up to 33,600bps. The letters "bis" attached to a standard indicate that it's version 2. The letters "ter" or "terbo" mean version 3.

In order to confuse you, there are error correction standards that have names which sound very much like the modem standards: The V.42 error-correction standard for modems specifies the use of the MNP4 protocol for V.22, V.22 bis, V.26 ter, and V.32 bis modems.

So, with V.42 bis compression, data is compressed at a ratio of about 3.5 to 1, which can yield file-transfer speeds of up to 56,000 bps on a 14,400-bps modem. Got that?

Higher throughput data lines have been available for some time. They include switch-56 and frame relay. As the name implies, Switch 56 is switched, which means you dial just like you would on a phone line. Frame relay is the standard for a low-overhead packet-switching protocol that provides dynamic bandwidth allocation at speeds of up to 2Mbps but can be consumed in increments of 56k. Frame relay is rapidly becoming the medium- to high-speed access solution of choice in the United States. It's not switched (though it can be routed), so it is "always on."

Both of these data lines offer 56,000 bits per second service much more reliably than POTS. With compression, this is pretty quick. Pricing for these types of circuits has fallen over the past few years. For instance, a frame relay connection at this time in Tucson is $72.03 per month.

ISDN (Integrated Services Digital Network, or "I Still Don't Need it" for some cynics) is a ubiquitous service in Europe and has become much more so in the U.S. as local telcos update their switching gear. ISDN is a digital phone line that provides 128,000 bits per second service on a dial up line. It offers a separate small signaling band or circuit so that connection times are almost instantaneous, and many added services such as caller ID, call forwarding, voice mail, and X.25 data signaling are built into the phone and are essentially idiot proof and transparent to the user. The fees for this type of service are falling and vary from $20-80/month across the country. You can hang your Fax, phones, teleradiology stuff, and computer on your home system and use the bandwidth dynamically to do many things simultaneously. The only bad news is that your current telephone handsets won't work with ISDN and many phone providers aren't very good at deploying it.

ISDN comes with some additional acronyms, which are helpful to throw around. Basic Rate ISDN (BRI) is an ISDN service that offers two "bearer" channels (B) with 64Kbps bandwidth that can be used for bulk data transfer plus a "data link" (D) 16Kbps channel for control and signaling information. Thus, it is way cool to ask, "Does that BRI offer 'two Bs and a D'?"

PRI (for primary rate interface) is 24 ISDN lines bound together as T1 service, offering 1.5 megabits per second. In PACS time that's 1.5 seconds per non-destructively compressed 2Kx2K image. Pretty swift. This is usually configured as a non-switched (always open) service as well, and bandwidth can be allocated or shared amongst competing service demands. The price of this service, which used to be distance-based and leased between fixed points only, has also changed for the better. You can now get T1 service directly to your telco for about $300-400/month. Past the telco, long distance charges apply, but if you're moving images over the WAN in a city, this is still pretty cheap stuff. (Nerd Stuff: In Europe, E1 is the equivalent of a North American T1, and it operates at 2.048Mbps. In the U.S., T1 is the basic carrier and operates at 1.544Mbps. The difference is that an E1 is 30 56K lines while a T1 is 24 56K lines.)

Then there are DS-3 (T3) lines. A DS-3 circuit carries in one multiplexed signal stream the equivalent of 28 T1 circuits. It provides 44.736Mbps of bandwidth. Switched Multi-Megabit Data Service (SMDS) is a high-speed metropolitan area network service for use over T1 and T3 lines. SMDS' deployment seems to be stalled by the enthusiasm for Asynchronous Transfer Mode, although SMDS is not in and of itself incompatible with ATM. More recently, fiber backbones are bundling together multiple OC-3s (155Mbps) into OC-12 (622 Mbps), -24 (1.24 Gbps), -48 (2.5 Gbps) and even -192 (10 Gbps) for some truly jaw-dropping bandwidth!

And then there's the Holy Grail: ATM. Now pay attention here, because ATM may make it to your desktop from the United Arab Emirates. ATM is a method of data transmission used by Broadband ISDN. Speeds of up to multiple gigabits per second are possible, and it is capable of carrying voice, video, and data. ATM has been embraced by the LAN and WAN industries as the solution to integrating disparate networks across a large geographic distance. It has been called cell relay - not to be confused with frame relay.

2. Please Ignore the Network Layers

A network can be described at seven layers, often referred to in the business as the OSI network layer model. The seven layers are:

1. Physical (for instance 10BaseT Ethernet or Token Ring). 2. Media Access (a specific driver for a network card). 3. Network (e.g. I/P Internet Protocol, SLIP, PPP, or X.25). 4. Transport (TCP - Transmission Control Protocol). 5. Session (e.g. Telnet, FTP, SNMP, or SMTP). 6. Presentation (e.g. Remote File Service) 7.Application (e.g. NFS - Network File System).

Unfortunately, a detailed discussion of these layers will be (a) confusing, (b) unhelpful, and (c) impossible because I can't figure it out either. In reality, this network layer scheme is an abstraction into which a reality variable conforms. Stacks of middleware are responsible for interfacing the network adapters to the software. It is easier to start with what works and what is out there, and then learn the layer model in retrospect. Additionally, knowing a few simple acronyms will allow you to get through this without knowing very much about the layer model at all. (Trust me, they've worked for me for years.)

NetBIOS is a protocol developed by IBM that governs data exchange and network access. Because NetBIOS lacks some layers, in particular a network layer, it cannot be routed in a network. This makes it difficult to build large internetworks based on NetBios. Examples of NetBIOS-based network operating systems include IBM LAN Server and Artisoft LANtastic.

NetBEUI (pronounced net-booey) is Microsoft's infinitely more pronounceable version of NetBIOS. NetBEUI stands for NetBIOS Extended User Interface. It too is a protocol that governs data exchange and network access. Again, because NetBEUI lacks a network layer, it cannot be routed in a network. This makes building large internetworks of NetBEUI-based networks difficult, but at least Bill Gates makes money. Windows for Workgroups uses NetBEUI.

IPX/SPX. IPX (Internet Packet eXchange) is a low-level network layer communication protocol used by Novell, which for the most part is supplemented with SPX. SPX is Sequenced Packet eXchange, a transport level for network data exchange that uses the IPX network layer protocol. (Since this is so circular, Bill Gates still makes money.)

TCP/IP. Transmission Control Protocol/Internet Protocol is the protocol used extensively on the Internet. It is the transport-level protocol. It is also widely used in corporate internetworks because of its superior design for WANs. TCP governs how packets are sequenced for transmission. The IP part of TCP/IP is the lower network layer. The term "TCP/IP" is often used to generically refer to the entire suite of related protocols. Conveniently, DICOM 3.0 uses TCP/IP as its communication protocol.

DICOM is a standard that is a framework for medical imaging communication. It is based upon the Open System Interconnect (OSI) reference model, which defines a 7-layer protocol. It is an application-level standard, which means it exists inside layer 7 (the uppermost layer). DICOM provides standardized formats for images, a common information model, application service definitions and protocols for communication.

SLIP, or Serial Line Internet Protocol, is used to run TCP/IP (Internet Protocol) over serial lines, like telephone lines. So, you will get a "SLIP" account from your local Internet provider, or if you prefer to pay CompuServe or America Online, they will give you a SLIP connection through your net launcher.

PPP, or Point-to-Point Protocol, provides router-to-router and host-to-network connections over asynchronous and synchronous connections. It is considered a second-generation Serial Line Interface Protocol (SLIP). Again, Internet providers will now offer a PPP connection option rather than a SLIP because it is more stable and less prone to interruptions and disconnections. Most people are using this now instead of SLIP. By the time you read this, FAX machines will have PPP protocols built into them. Next year, toasters too!

Teleradiology

Teleradiology is the process of sending radiology images from one point to another through digital, computer-assisted transmission, typically over standard telephone lines or a local area network (LAN). Through teleradiology, images can be sent to another part of the hospital, or around the world. Teleradiology was the first widely deployed implementation of telemedicine and PACS technology. It is considered a mature technology at this time and is firmly planted as a PC application in the marketplace.

A. Teleradiology Systems

There are three classes of teleradiology systems: (1) on-call, (2) off-site, and (3) in-hospital (mini-PACS). "On-call" systems are most frequently used for after-hour, "on-call" applications. Typical "on-call" teleradiology systems consist of a transmitting unit and one or more receiving units. Usually, both the transmitting and receiving units are based on standard personal computers.

"Off-site" or "overread" systems are used most frequently by radiology groups and hospitals to centralize readings and/or coverage, increase productivity, expand reading networks and consolidate subspecialty work. Off-site teleradiology systems are especially important to rural medical facilities.

"In-hospital" systems, or "mini-PACS," are used to move images electronically throughout one facility (i.e., from the radiology to ICU, ER, OR, and clinics) over a local area network (LAN).

B. Image Transmission

For digital modalities (CT, MR, ultrasound, nuclear medicine, computed radiography) images can be captured either by a video capture (frame grabber) board which connects directly to the composite video signal of the image processor or the image display console, or digitally by connecting directly to a modality to a workstation over a network (i.e., ethernet).

Either a video camera or a film scanner can digitize hard-copy plain films. Video cameras (commonly referred to as a "camera on a stick") were in fact the method for digitizing any image for transmission as recently as 5 years ago. Typical camera systems utilize a light box designed to illuminate radiographs, an extension arm for holding the camera above the film, and a high sensitivity video camera with zoom lens. This is inexpensive, but as you might expect, produces poor results.

Hard-copy plain film scanners/digitizers arrived on the teleradiology scene a few years ago. There are two basic types of film scanners: (1) laser digitizers and (2) CCD (charge coupled device) digitizers. There is considerable discussion about these technologies. Laser digitizers clearly provide better images while CCD digitizers seem to offer very good value for the money. Certainly, CCD technology has improved considerably over the past few years and today's high-quality CCD digitizers produce images that are as reasonable comparable as top-of-the-line laser digitizers within OD ranges under 3.0. An excellent explanation on the gory details on the difference between these two technologies can be studied on the Lumisys home page http://www.lumisys.com.

Depending on data transfer rate requirements and economic considerations, images can be transmitted by means of POTS, ISDN, switched-56, microwave, satellite, and T1 telecommunication links.

Today most teleradiology systems run over standard telephone lines. Over the next couple of years, we should see substantial migration to switched-56 and ISDN (Integrated Services Digital

Network) lines, which offer higher speed and better line quality than standard dial-up phone lines. Other high-speed lines, including T1, cable modems, and xDSL, will also become more popular as prices continue to drop.

"How fast can you transmit an image?" is probably the most frequently asked question of teleradiology sales reps. - the answer is not as straightforward as you may think. An "image" is what is shown on the display monitor; it could be a single CT slice or an entire 14"x17" film. If a vendor tells you it will take 15 seconds, he isn't really answering the question.

Image transmission time is directly proportional to the file size of the digitized image. The greater the amount of digital information in an image (i.e., the larger the matrix and the larger the number of bits per pixel), the greater the amount of time required to transmit the image from one location to another.

The only way to decrease the transmission time is to increase the modem's speed or reduce the number of bits being sent (compress the image). The following formula is used to calculate the time to transmit an image.

Transmit time = matrix width x length x (matrix depth + 1 bit) x 1 x .8 compression ratio baud rate

The addition of "1" to the matrix depth is a housekeeping bit which is the "stop bit" that the modem adds to each byte. The multiplication by .8 reflects a realistic 80% efficiency of the modem at the rate you select.

Modem protocol overhead, turnaround time, and typical phone line interference will cause all modems to transmit more slowly than their published capacity (20-30%). When a phone line has noise, i.e. music from a local radio station, voices from your other line, a crossed line, or clicks caused by line switches, the modem will "hear" the noise, interpret it as bad data, and re-transmit. If the disruption reaches a high enough level, and it often does, the modem will reduce its speed or disconnect. The degree to which a modem slows down is called a fall-back rate. The fall-back rate is different for each manufacturer's modem and varies among phone calls. The modem drops from 28,800 bits per second to 26,800 bits per second and so on at the first sign of line interference. A smaller fall-back rate makes phone line interference less noticeable. Because fall-back rates make modem speed so variable, this formula and discussion assumes a completely clear line and modems transmitting at their top speed.

As you investigate teleradiology choices, you should pay attention to the fall-back rate of a system's modem. A large fall-back rate would be a source of frustration for you (some modems have a fall-back rate of 2,800 bits per second or even as much as half the modem's speed).

C. Teleradiology Guidelines

The American College of Radiology Teleradiology Guidelines were approved in November, 1994. The guidelines provide recommendations governing the use of teleradiology, defining such areas as teleradiology goals, personnel qualifications, equipment guidelines, licensing, quality control and communication. The following information represents highlights of the published guidelines.

Equipment guidelines cover two basic categories of teleradiology systems for primary interpretation: small and large matrix systems. Small matrix systems, covering CT, MRI, U/S, NM and digital fluorography, require a digitization system with a 0.5K x 0.5K x 8 bits array or better. Large matrix systems, covering digitized radiographic film and CR (computed radiography), require a digitization system with a 2K x 2K x 12 bits.

All teleradiology systems must include annotation, window/leveling, image sequence selection capabilities, and the ability to select compression ratios to improve transmission rates and reduce archiving/storage disk space.

A patient database, to include the patient's name, ID number, date, type of exam, and type of images, should be an integral part of all teleradiology systems at either the transmitting or receiving site. Additionally, transmitted images should include annotation of the matrix size, bit depth, compression, and type of image processing used (such as edge enhancement).

The ACR guidelines make specific recommendations for display stations used for final interpretation. One recommendation is that these stations should provide gray-scale monitor luminance of at least 50 foot-lamberts. It is also noted that small matrix display stations, used for image interpretation, must accurately reproduce the original study. Finally, large matrix display stations should include interactive window/leveling and magnification features, rotating, flipping and image inversion features, and the ability to make precise linear measurements.

Because of confidentiality of patient images and data, the guidelines state that teleradiology systems should include network and software security safety features. A quality control program should also be in place to assure optimum performance of teleradiology equipment.

While most teleradiology systems purchased over the last decade were for on-call purposes, the past two years have seen a rapid increase in the use of teleradiology to link hospitals and affiliated satellite facilities, other primary hospitals, and imaging centers. A number of the enabling technologies needed for effective over-read networks, such as more affordable high-speed telecommunications networks and improved data compression techniques, have matured in recent years.

Input/Output Devices

And now for something almost practical: a discussion of the devices you're used to seeing, buying and selling in a radiology department on a daily basis: laser cameras, spoolers, frame grabbers, etc. Back when multiformat cameras were high tech, it was enough to simply photograph a copy of the CRT display onto a sheet of film. You had several choices: 4 on 1, 6 on 1 or 12 on 1. Now no one buys a device unless you can peel shrimp with it. And for good reason.

A. Laser Camera Spoolers

Most manufacturers only build in a fixed number of "ports" onto their laser cameras, limiting your department to that pre-determined number of inputs for that camera. You may be the lucky owner of such a device. If you are, don't fret you can find a way to expand in the future: Enter the laser spooler. Here's a piece of info worth knowing: Spooler stands for Simultaneous Peripheral Operations On Line. In keeping with the newer trends in computing, distributed computing is cool. You can now buy a laser spooler, previously running only on "workstations" but now running on a "lowly" PC which will accept multiple (usually multiples of 4 inputs) per frame grabber board and accept as many images from as many cases as you throw it. The spooler's job is to group the images into cases (i.e. folders) for film printing. If you prefer, you can edit the images or even annotate the images in a case before printing. Need more inputs? - no problem The PC works on a frame grabber card (see below). Most frame grabbers in the medical image world support multiple inputs (often 4). If you need to feed more devices in than one frame grabber card permits, just have your PC guru open up the box, seat an additional frame grabber card and configure the software.

With the advent of a TCP/IP network protocols like DICOM, newer laser printers will be "networked". That is, they will have a network interface card, and have a home run of wire back to the local hub. Then any other scanner or device on the Ethernet LAN can use the printer to print or archive to archive. In essence a virtual connection is established through the LAN between the capture device and the printer, through the hub. So, for any DICOM CT scanner, MR, ultrasound, angio device, or even a digital fluoro room, you'll be able to address the printer or a print spooler, archive and review station(s) all at once! It should be just as easy to address these devices over the WAN and print from a remote site. Instead of a hub, telecommunications through a router will be required, and throughput is subject to all those telecom issues discussed above.) Again, depending on the image traffic flow, you may still require a spooler for a given camera. But all in all pretty smart, huh?

Between today and that arrangement however is, for most of us, a long journey. Many devices are not, nor will they ever be DICOM compliant. To obtain images from these devices in digital format or put them on a DICOM network will require one of several possible secondary input devices or black boxes. The secondary input devices include laser film digitizers, video frame grabbers and CCD film digitizers.

B. Frame Grabbers

Video frame grabbing basically consists of digitizing a frozen frame from a video source. Since your video display is designed for your eyes, and your eye basically sees 256 shades of gray (or 8 bits - that's 28), frame grabbing is an 8-bit process (unlike DICOM which is 16 bits per pixel). All ultrasound and some Nuclear Medicine devices output "standard video", or the same synchronized signal which comes out of the back of your TV and goes into your VCR. This is affectionately code named RS-170 (or PAL in Europe).

Everything else, CT, MR, angio, fluoro, etc. puts out (you guessed it) a non-standard video signal, proprietary to the manufacturer. The number of scan lines and frequency varies enormously so that your standard frame grabber boards from PC magazine will only work on the minority of medical devices. These non-standard specifications require boards that can be adjusted or tuned to the particular video signal from which you desire to capture. Since the market is small, gray scale is important, the images are large in terms of pixel size, and the consumers are picky, the price is stiff. Depending on the signal a good grabber board will cost in the neighborhood of $5,000. These boards are used in a variety of locations ranging from teleradiology systems, to laser spoolers, to laser film printers.

C. DIGITIZERS: Lasers and CCDs

Film scanners are used to digitize films, e.g., convert the analog information stored on the x-ray film to a digital data stream. Film scanners feed a digital data stream to a PC through one of the PC's I/O (Input/Output) data ports. The light source used in Laser film scanners is a finely focused, coherent and monochromatic infrared or visible red laser. The Laser beam must be precisely positioned and scanned through the film under computer control, while the light transmitted through the film is efficiently detected by high gain, very low noise devices and amplifiers then applied to digital conversion electronic circuitry for input to the host PC. This makes the device extremely precise, more costly than a CCD based scanner (see below), and requires certain component features, which make it somewhat physically heavier than CCD scanners It is important to understand that there are, however, very technical and important considerations which usually make the properly designed laser scanner's optical density steps scanned narrower, more linear, and with its controlled spot of scanning laser light, yields very accurate spatial resolution independent of film density, therefore digitizing, without distortion, the information on the x-ray film. Hence, a laser scanner is more favorable for primary medical imaging applications than CCDs. (The technical advantages and characteristics of laser and CCD scanners are best studied and gleaned from a white paper on the web site of Lumisys). http://www.lumisys.com.

Also, the best laser scanners approximate a scanning densitometer, where the digitized information derived from the scanner accurately represents the actual film density information on the film at each and every point on the film. This is an achievable characteristic pertaining to the point source of light and special light processing techniques in a Laser scanner transilluminating the film compared to the CCD scanner with a line of white light derived from a fluorescent tube as a light source.

CCD film scanners are physically quite similar to document scanners available in the commodity marketplace. A CCD is charged coupled device (an array of solid state photo detectors), similar to that contained in your home video camera or camcorder, and is used to capture and digitize images. The digital data can be, for example, fed to a PC through a SCSI cable and board. Unlike scanners in the commodity markets however, films are generally larger than paper (up to 14"x17"), they are transparencies, they are gray scale, and the users, physicians, are picky. Ergo price goes up.

CCD scanners, manufactured under cost performance constraints, cannot achieve the standards of a laser scanner due to the inherent electrical and thermal noise, and non-linearities associated with semiconductors. Still CCDs are reasonably priced these days at about the $17,000. The current vintage CCD scanners can provide 12 bits of gray scale data over a limited dynamic range (the last 4 bits in the 16 bits are usually 0's anyway) and will scan up to 4kx5k.

Whether or not the file output of these scanning devices is DICOM is really up to the software being used, not the scanner manufacturer. All the scanner does is transilluminate the film and provide a digital data stream of 1's and 0's. Once the PC gets that stream its up to the software manufacturer to package the image in DICOM format. Some do, some don't.

D. CR

Computed Radiography (or CR) has been around for about 20 years. The technology basically consists of substituting a storage phosphor plate for a light-emitting phosphorescent plate within the radiographic cassette. While this may seem painful, you might recall that a light-emitting phosphor plate converts the x-rays that make it through the patient on exposure into light; the light in turn exposes film, and the film is developed with chemicals just like it is by your corner drug store. Granted, its all black and white film, big sheets not 35 mm., but the process is basically the same.

In CR, although the phosphor in the plate emits a small amount of light, the technology is designed to basically trap the x-ray energy which make it through the patient as a latent image. The plate is then passed under a laser light beam, which excites the trapped energy and releases it in the form of visible light at that time – i.e. at a time after the exposure. This latent image is then ‘read’ out by a series of electronic and amplification devices, which in turn make the ‘image’. This device is called the ‘reader’. Once the electronic image exists it can be mathematically processed, filtered, and tinkered with to make it better. Then it can be displayed on a monitor, or if you sell film for a living, printed on good old fashioned film.

Patented and commercialized by Fuji, originally the benefits of CR (the output of which was originally only printed to film) were primarily the reduction in repeat/reject rates for portable x-ray films, yielding more consistent and desirable film densities, and the ability to subject images to unsharp masking and other algorithms prior to printing them to film. The financial effects of saving a reject/repeat rate, which in various departments will range from 3-7%, is of course not terribly great. More recently, however, the ability to leave the images as soft copy and move them around the healthcare facility for simultaneous use has became more feasible with the adoption of the DICOM standard.

Even so, almost 80% of facilities which own CR devices still print to film for reading or archiving. A similar percentage planning on buying CR in the 1997-99 time frame also indicate that they too intend to print the images once acquired. Clearly the financial benefit of a CR device, ranging in price from $125,000-$250,000, and still attached to a film printer and processor is dubious if all that is achieved is a maximum 7% reduction in rejects and repeats. In fact, a recent study at the University of Pennsylvania put the net present value of a CR study at negative $7/exam because of the high capital costs of the reader.

Softcopy may finally deliver a financial benefit through DICOM connectivity despite nearly 20 years of sluggish market penetration if printing is abandoned for softcopy distribution and interpretation. Also offering to modify the financial equation is the entrance of other companies into this arena, including my one close to home: Lumisys. We fervently believe that a device which changes the financial equation favorably will necessarily favor the use of softcopy, rather than film in a DICOM environment, distributing very low-cost CR devices into low volume and often remote locations. Native DICOM digital images everywhere is the battle cry. Whether this strategy is better than film or in your future remains to be told.

On the footsteps behind softcopy CR is the archival problem. People began to ask, "Why print the film just to file it?" As people began to archive images on magnetic media, they began to realize that using the DICOM 3.0 format would be a swell idea. So, everyone seems to be scurrying in order to get CR image files into DICOM 3.0 format, primarily for the purpose of archiving and query/retrieval.

E. The Brave New World of DICOM

When DICOM was in the womb it was known as ACR/NEMA. ACR/NEMA Version 2.0 specified, amongst other things, a special 50 pin parallel port connector through which one could attach devices like CT scanners to a networked environment. One of the major problems with this approach was that it was a relatively unique plug, which required manufacturers to engineer and conform to a hardware and software standard promulgated by a standards organization with unproved market power or merit. It was also a time of fabulous hardware sales with increasing price competition. Manufacturers looked to the marketplace to underwrite this engineering and hardware change, usually asking about $25,000 for such an output device on a scanner.

Few buyers bought them. But with the advent of DICOM 3.0, in which well supported, widely deployed communication protocols (TCP/IP) are implemented over a variety of relatively ubiquitous media and network infrastructures (e.g. Ethernet, ATM, fiber, UTP5 wiring), manufacturers seem more interested. PACS systems, which probably represent the greatest growth opportunity in radiology departments during the last half of this decade, are sold to a health care sector, which is now very well, funded (information systems). PACS in effect require this type of connectivity. Without it, you'd have to buy all of your hardware from a single vendor - a situation, which is a) unlikely and b), ignores the variety in the existing installed base.

As a result, most of the major manufacturers have embraced DICOM 3.0. Take Home Message: In 1997, buying a device without DICOM 3.0 capabilities or guaranteed after-market delivery of the same should probably be considered managerial malpractice.

F. Go To the head of the DICOM Class

A DICOM device can be anything from a reading station, to a preliminary review station (for instance in the ICU), a CT or MR scanner, a computed radiography (CR) laser reader, a tape archive, or a DICOM gateway. But the label DICOM alone doesn't tell you which provider classes the device supports. In order to understand the interconnectability of DICOM devices for your intended use(s) you will need to become familiar with a few highlights of the seven volumeDICOM standard. This will help you make sure you get what you need when it comes to buying DICOM devices.

A DICOM "storage class provider" (SCP if you speak hip DICOM) device can provide images in DICOM file format, over a network with a DICOM compliant protocol (TCP/IP). That is, it can "push" images to things like a DICOM compliant laser printer, which is a "storage class user" (or SCU), DICOM SCU spooler, DICOM SCU "workstations" or DICOM SCU archive.
That doesn't mean you can query the provider class device, it just means the device can "push" DICOM images to it and it will accept them. A typical provider class device might be a CT scanner or workstation.

A DICOM query/retrieve provider is any DICOM device which can query other DICOM query/retrieve devices on the network to find cases and images and, if desired, retrieve them; conversely, it can provide the results of a DICOM query originating from another DICOM query/retrieve device and, if requested by that device, deliver the requested DICOM image files to it over the network.

Now you might imagine that being a DICOM storage provider class only device is kind of dumb. It’s like someone who only knows sign language and is blind. She can talk to you, but you can't talk to her. She can push images to other devices, but she can't really communicate with them (i.e. ask it questions). This certainly is not what you'd pick these days for say a CT scanner, but if you have an older (non-DICOM) scanner this is about as good as you'll be able to do by inserting a "black box" between the scanner and the network, so as to make your scanner "DICOM compliant".

A storage class device without query/retrieve capabilities is like a black hole. You can put DICOM images in, but you can't get DICOM out. In all fairness, some manufacturers created DICOM storage class devices as an additional service to their prior generation proprietary storage devices and/or ACR/NEMA 2.0 storage devices. The engineering in query/retrieve isn't trivial. So the product roll-out has generally been to provide the storage class service to enable you to hook up the new DICOM acquisition devices; following this query/retrieve will be provided so that once you add DICOM workstations, you will be able to get DICOM out when you want it.

DICOM workstations of course need to a) display images pushed to them by acquisition devices ("provider class"), b) store images pushed to them for later display ("storage class"), and c) be able to find and retrieve images on the network from storage and acquisition devices ("query/retrieve class"). While you can get along for a while without query/retrieve, you'll want it once you start doing a substantial amount of image surfing at any workstation.

G. DICOM Wannabees

As you wire your imaging empire you will sadly look upon that old CT or MR scanner and lament that you cannot afford to replace it, but you'd sure like it to be on the network so you can print, archive and read at your workstation from it. Enter the lowly black box. There are a few manufacturers who make such devices, and they do not come very cheaply. (But what do you expect, the market is small and the market self-limited) If you need one, buy it and don't whine. But be aware that you're trying to get an extra few hundred thousand miles out of a Volkswagen bug. You won't get query/retrieve and certain portions of the DICOM image header won't be provided, most notably organ information. But you will be able to print, archive and read from your workstation (that was the original purpose wasn't it?).

There are always a few ways to skin a cat, and black boxes are no exception. Each box will have its own advantages and disadvantages. Additionally, other solutions can provide the same output and gain you additional functionality. For instance, you can use a PC with a frame grabber in order to put images into (albeit an 8 bit) DICOM format, and at the same time use the same PC as a laser spooler for your department and/or a teleradiology capture/send station during the wee hours.

Other middleware devices will bridge the gap between your existing gear and the brave new DICOM world. Those gotta-have-it-first guys who bought early ACR/NEMA 2.0 compliant PACS systems, and you fools of the latter days who continue to buy non-DICOM compliant but networked imaging devices will be faced with mixing DICOM and non-DICOM images on your network. You want that ACR/NEMA 2.0 thing from vendor 1, your old CT scanner with a black box from vendors 2 & 3, and your brand new DICOM Swiss army knife from vendor 4 to talk Lebanese? Well, in these circumstances you will probably need a gateway. This is PACster talk for middleware, the price of which will always exceed your concept of what is fair. These are growing pains. And to get files, applications, computers and medical devices of a different vintage to inter-operate is difficult and expensive. Add your intolerance for downtime and the FDA's regulatory interests on behalf of our patients and it's downright expensive. So don't whine. These aren't exciting products or easy sales for PACS vendors either. A number of PC, mini-PACS, and major manufacturers are offering such DICOM gateways. While they may seem expensive now, they will ultimately be seen as inexpensive methods by which to network older and non-standard devices without replacing them.

Basically this is what all of us want: DICOM-in/DICOM-out or (DI/DO if you like acronyms). In the process of getting DI/DO you will likely want to interface your secondary capture devices (e.g. your teleradiology gear) and your older scanners to a DICOM network. To do this, a variety of proprietary and older ACR/NEMA file formats will need to be translated into DICOM 3.0 format and thereafter provided to other DICOM devices (e.g. workstations) after issuing a DICOM query/retrieve request. Such devices are called gateways.

A gateway can be a standalone computer or a gateway service can be provided from inside an existing computer on the DICOM network as an added "service".

Because image files are large and file format conversions are pretty computing intensive, the hardware requirements for a gateway or gateway service are substantial. At the PC level a multiprocessor server or at least a Pentium computer with server level hardware configuration is necessary. On the "service in a DICOM box" side, a Sparc 10 or so will do. Still, you must realize that as your network DICOM traffic increases, your gateway service will become increasingly important, even mission critical. Having the gateway as a service on a DICOM workstation or archive device used for other purposes may make the gateway a) unreasonably slow, b) may deteriorate the other services provided on that workstation or archive, and/or c) may create unbearable failure if two services crash simultaneously.

All in all, gateway "service" may be a reasonably priced way to start with a low volume DICOM gateway, but you must realize that you will likely be locked in to the same vendor when you need to upgrade to a stand-alone gateway server.

H. Archives

Archive devices will vary substantially in price, size, and ability to integrate and scale to your needs. Some archival media are CD ROMs, optical disks, RAIDs and digital tape. I hesitate to get involved in a directed discussion of choosing between devices in a Primer. I will however offer some free advice (be warned, you usually get what you pay for).

First, determine where you are going (i.e. what your electronic empire will look like when you are finished) in order to take your first step without bungling. Most of us have a different destination and a generic discussion just won't do. If you are an imaging center or a 100 bed rural hospital and never expect to have to integrate an electronic archive with another bigger enterprise, you can probably safely choose some simple approach like CD-ROM or DVD in the next few years. If you think that you ***might*** have to integrate to an electronic medical record, a bigger hospital, an integrated delivery network, then planning and flexibility is critical. You are making a decision of a different order than buying your next fluoro room or CT scanner: don't kid yourself. My advice is this: a) if you need help, get a good consultant (there are several in our industry and I will be happy to refer you to them if you need them); b) make sure the purchase includes state of the art, mission-critical hardware from a reputable manufacturer (you may buy it through a System Integrator, but make sure you know where the hardware comes from); c) buy your archive from someone who knows archiving, not fluoro rooms. That doesn't necessarily mean the biggest name, or the biggest company. It means, buy a reliable archival *solution* from a firm that will be there with capable people whenever there's a problem.

You are buying a *solution* not a simple device. Whether its from GE, Siemens, Agfa, your MicroAge dealer, Mantis, or your regional teleradiology and PACS vendor, anyone who sells you a DICOM archive is selling you software which is at least in part sourced from third parties (e.g. Informix, Oracle, Lumisys, etc.), hardware which is sourced from third parties (e.g. StorageTek, Disc, etc.), and integration, training, support and maintenance *services*. This is a solution. It works like any imaging chain: the weakest link in the chain determines its overall strength, period. It doesn't matter how big the company is, if the local service and support is poor, it will be like you padlocked your file room and lost the key. You will be closed for business when your archive is down.

Advanced Networking Stuff

In Chapter III we discussed basic networking. In order to get a true feel for where your PACS network can take you, it is necessary to understand some more advanced devices and concepts. If you don't understand significant portions of this chapter, don't fret it only took me 3 years as CEO of a PC teleradiology/networking company to understand and learn them. You will come to understand these things over time, or through experience, or by re-reading this until it makes sense.

A. Bridges, Routers, and Hubs

As you know, a hub is a device which interconnects various nodes to one another on the LAN. In essence it's a dynamic switch. A bridge connects two networks of the same access method, for example, Ethernet to Ethernet or Token Ring to Token Ring. A bridge works at the OSI's Media Access layer, and is transparent to upper-layer devices and protocols. Bridges operate by filtering packets according to their destination addresses. Most bridges automatically learn where these addresses are located, and are thus called learning bridges. A bridge is a device, which takes a communications line coming in or going out and does the same thing.

That is, if you have an ISDN phone line coming in to your house or office, hook it up to a bridge and then the bridge to your computer, your home or office computer will act just as if its on the LAN that its connected to over the ISDN line. Now you can do this with a plain old phone line or a T-1 line. The throughput will be faster or slower depending on the bandwidth of the telecommunications, but the concept is the same.

Bridges are computers or devices that interconnect LANs. Protocols that don't support inter-networking (like NetBIOS/NetBEUI) are bridged. Router interconnects LANs using protocol dependent routing information. They are typically used with protocols such as TCP/IP, IPX/SPX.

That's when life was simple. Now products overlap, so there are bridge/routers, multi-protocol routers, and gateways. Bridge/routers are single devices that combine the function of bridges and routers and have become quite popular. Routers may be multi-protocol routers, which support different combinations of network layer protocols. Finally, there are gateways. Gateways are devices that do the brute force of translating between network protocols.

Despite their popularity on Wall Street, these devices are Band-Aids. When the widespread use of multimedia starts to clog the networks in our environments, these devices will become increasingly expensive Band-Aids. Routers will become much less capable of handling traffic flow - delays in traffic flow which are not acceptable to real-time interactive data such as voice and video.

B. RAIDS

RAIDs (short for Redundant Array of Inexpensive Disks) involve disk mirroring in which data is written to more than one hard drive. This offers failure protection and also, depending on the scheme used, will improve disk retrieval performance critical to overall PACS performance. Five classes of RAIDs are available; the details of which are beyond the scope of this monograph.

C. Client/Server

Client/server computing is the new rage. If you offer it, you are way cool. But most people don't know what it means. In the old days a mainframe computer housed both the data and the applications. Terminals were used and literally called VTs or video terminals. In effect they were dumb clients. They had no brain, hard disk etc.

As PCs and workstations became abundant, applications were located on the stand alone PC. Then came networking which allowed PCs to access each other's files. This was known as file sharing.

Although there were problems with file locking, this actually will get many small offices through their needs. With regard to file sharing, without a network operating system it is clear that only one copy of a file could be open at anyone time; otherwise, two users making changes to the same file would corrupt it. Indeed, this is a major function of a NOS (network operating system).

In the case of databases, many users even through file sharing can open the same database. Locking can occur at many levels. The most popular methods are record locking and field locking. Record locking prohibits the second user access to modify a particular record; whereas field locking allows access and changes, narrowing the locking down to one field in a given record.

Indeed this type of file sharing or "groupware" can be had with almost any Windows application. Put the application on one machine and you can run it simultaneously from many others (you will however need enough licenses from the software firm to do this legally). Put the application and the data on a dedicated file and/or application server and you will even further improve application and network performance. First, no one PC is slowed by other PCs accessing it for the application. Second, the housekeeping chores of back-ups, disk defragmentation, and uninterrupted power supply are centralized and therefore simplified. Third, adequate resources of RAM, disk space and CPU speed and power can be more cost-efficiently provided at the central source.

The next network architectural change came in the form of the client/server concept in which the data is separated from the application and may conform to a variety of (mostly database) access standards so as to allow access of the same data by multi-vendor applications. This has permitted organizations to buy XYZ company SQL database and run a number of front-end applications on it. The application accesses the data via a variety of protocols and/or queries and presents the results to the end user through the application.

The data remains separate from the application; the application can reside on a file server or at the local workstation. In fact there may be multiple different applications running on the same network, addressing the same database. Conceptually, the "server" is really a service or a proxy agent. It is the middleman or policeman, which facilitates data access, data changes, and collisions. The server is an agent, which keeps event logs. The server may thus "replicate" itself to other servers and, in so doing, widely distribute the network.

Indeed, the concept for distributed computing and client/server computing are intimately interwoven. These services and databases may be replicated over LANs and WANs to better distribute resources for the network's demands. Multiple, redundant points of failure are created as well, building better fault tolerance into the system.

Client-server computing has numerous advantages for the medical world and for the field of radiology in particular. It permits workstations to achieve computing power previously only available from mainframes -- at a fraction of mainframe costs. By efficiently dividing resources, client-server computing reduces network traffic and improves response time. This offers a significant advantage to physicians who need to receive images quickly and who require real-time image navigation and manipulation to perform diagnostic tasks effectively.

Client-server computing facilitates the use of GUIs and OOUIs making teleradiology and PACS applications easier to use and more responsive. Clients and servers can run on different platforms, allowing end users to free themselves from particular proprietary architectures.

While PACS systems are not universally yet designed this way, there is little doubt that as they are deployed client/server architecture will become required. If for no other reason than the large image file sizes and the widely distributed demand for those images through the health care network will be made over network media of variable competence to sustain the necessary data throughput and still maintain network functionality.

In a similar way, the mainframe construction of Hospital Information Systems (HIS) and Radiology Information Systems (RIS) will begin to melt as full implementation of DICOM and HL-7 networking standards take hold and permeate growing healthcare enterprises.

Devices will be reduced to their simplest functionalities. CT scanners will scan, but the images will be addressed in an intelligent way to the tiered archive immediately, to the reading station where it will be interpreted and to the patient care area from where it was ordered. You won't buy a short-term archive at the CT scanner, one at the MR scanner, one as a miniPACs in ultrasound and Nuclear Medicine when you have those things in a more robust embodiment down the Ethernet.

D. Workflow, Mini-PACS, and "Scale"

Workflow consists of routing information in a predictable way, which mimics the ordinary flow of paper, film, voice and e-mail. Once the electronic infrastructure is in place to permit it, the elimination of paper and film only increases efficiency if the workflow is automated, and can be modified on an ad hoc basis, thereby reducing the human input necessary to move the (now electronic) elements. The banking and manufacturing industries have been doing this for a few years now and it is no longer a black art.

Workflow permits conditional routing of these elements (paper, images, voice and e-mail) as well as concurrent use. A variety of development environments and development tools, most of them object oriented, are widely available for commercial development and use. It is likely that healthcare information companies and even hospital/integrated healthcare system IS departments will develop workflow routing for their unique workplace needs once PACS and paperless medical records become significantly deployed.

It is essential that you, the purchaser, seller or consumer of PACS understand this concept before you buy your first piece of PACS equipment. The reason for this is one of scale.

When you look at a hypothetical miniPACS system for CT or MR from XYZ company you will think like this: If I can cut down on my film, my tech's visits to the darkroom (time) and archive to CD ROM, I can get an ROI in 12-18 months. This is a slam-dunk. And indeed it is.

But in three years, when you have 300-400 CDROM platters and a jukebox that holds 100 platters, getting Mrs. Jones' prior lumbar spine CT or MR for comparison may take a substantial amount of clerical time. This too, you may or will reason, is still a bargain, mostly because the demand will be low. But you must remember as well, the demand may not be local (i.e. in the radiology department), it may come from across town at the neuro or orthopedic surgeon's office (who has by then been attached to the hospital network with frame relay or ISDN).

If too many non-local sources are making random access demands on your miniPACS archive, you will have one grande headache in your department. Obviously, once this happens, it may be best to have your data on a tiered archive, which includes digital tape and optical disks. This is what we call scale. Sooner or later, given the recording media, you will have too much data to keep it all on-line and the nature of the media is not conducive to scaling if demands greatly amplify. It is true that you can be induced to add more CD jukeboxes and then increase the sophistication of the database software so that accessing the older studies is transparent to the client.

But, in effect, you are just delaying the inevitable and trying to make a slingshot into an assault rifle. CD ROM will not scale well in a hospital setting. It may do just fine in a stand-alone imaging center, but then again most of these are affiliated with groups, integrating health enterprises. They too wish to be on line.

Ultimately, such a Mini-PACS will have to be re-arranged to become part of a real PACS. It may be cheaper to pay for the manpower to transfer and re-index the data from the entire CD ROM archive onto the digital tape tank once one is acquired, rather than have a CD ROM disk jockey on your staff (but that wasn't part of the ROI plan was it?) So, treat these ROIs as only guideposts. Don't be surprised when the velocity of digital information demand within a short time far exceeds that of film in your current environment. In that setting the ROI doesn't reflect either what your current costs would be, nor does it take into account the growing demands on your personnel and PACS for information to sources outside your department.

Beyond this the ROI usually doesn't take into account down time - an inescapable feature of computers. Nor will it usually reflect system administration costs and time - costs, which will increase as increasing external demand for images increase. Nor does it reflect support costs - usually one doesn't pay a support contract on the file room. Nor does it reflect learning time for technologists, system administrators, and physicians every time you change archiving media. So… when you look at these ROIs, realize that there are hidden, incalculable costs which are totally unforeseeable. Enough said.

Archives

Very few topics are more boring than archives. It's kind of like trying to breathe life into a discussion about filing systems. But people do somehow get religiously impassioned about their views. As a result, people usually push their own solution, media, and/or device as the be all and end all. Building archival solutions is a huge undertaking. To get it right and make it work takes lots of time and money. So when someone is finished and has a product to sell you, they can't go back and re-do it because USA Today is pushing double-sided, quad density blue-light DVD at the annual Consumer Electronics Show.

Which brings me to the great deception: Media and devices are coming out at a furious pace. The prices are falling and terrific for consumers. But you shouldn't get confused between a mission critical function like a PACS archive with hierarchical storage management and your home DVD player. At home you want to watch Terminator 8, at work we take care of patients. Fidelity, reliability, mean times between failures, access times, prefetching, data migration, redundancy and similar issues aren't real important to your kid, Blockbuster video, or TCI cable. So listen carefully…

Departments will produce from half a terabyte to 3 or 4 terabytes a year, depending on their size and whether or not plain films are digital in the form of CR, DR, and whether you include Cath labs studies and studies from echocardiography labs. (This doesn't include the digital mammography issue coming in the next century.) No matter what, even a 100 bed, small department with no CR, DR, cath lab, etc. will produce a couple of terabytes over the expected 5 to 7 year statutory period of film guardianship for a radiology department.

So just think about the ho-hum, middle of the road rate of 1 Terabyte per year: That's a thousand gigabytes per year (one terabyte) or 85 gigabytes per month. Now before you get to tape, MOD, DVD or CD in that environment, most people would want on line access to a month to three months of images pretty quickly. If you look at the number of times studies get pulled from a file room, its pretty intense for the first couple weeks, moderately intense for the first 6 months, and then suffers a pretty precipitous decline after that. That's why most medium to large departments intuitively keep 3 or 6 months of jackets in immediate reach (on site), and have an 'active' file of in-patients subdivided from that. Form follows function.

By immediate access, most of us mean that once we've located a jacket or an electronic record we'd like to see the image or put the image up on a view-box in a few seconds. Maybe that's a leisurely walk to the viewbox, which takes 5-15 seconds; or an interminable wait at a CRT for 2-8 seconds. But it isn't a minute or two - it's in the seconds range. So, a CD jukebox (or DVD jukebox) is probably out of the question for immediate on-line access, unless it comes with some pretty beefy pre-fetching integration with the RIS and HIS to get the stuff onto a hard drive ***before*** we ask for it. (Hint: That's kind of like asking for a formula one engine inside of a VW bug - they just don't make 'em that way.)

So, the Fortune 1000 solution is to have a robot for MOD, DAT, DLT (discussed above and below), for 'near-line' storage and a RAID for 'immediate access'. That is, you have a workstation (local storage), a server with a RAID for fast hard-disk access, and near-line storage usually in the form of a robot. The size and type of storage you decide on probably is not terribly important if you can ***change the size and easily migrate the data*** in the future. That is, today you might choose DAT or MOD. In two years that will look a bit antiquated since newer sales may in fact be using holographic media (current research indicates that you can pack a gigabyte in a space smaller than a sugar cube).

For unknown reasons, people are always consumed by the hardware choices and so focus on them first and usually to the detriment of the real issue: software. So I'll walk you through the hardware first, but only because you won't read about the software if I don't do this first.
First, focus on getting very fast and highly reliable RAIDs so that the immediate demands on images can be satisfactorily met even when lots of people are contending for the same resource; and the RAID level should be reliable enough so that you can 'hot swap' disks that fail without any down time. In this business, up-time, hardware response times, and software reliability are everything. Inexpensive media is great, but if it doesn't deliver the images reliably and quickly, it's just cheap. If you buy great RAIDs that can be easily expanded and hot swapped, the front line troops will be happy.

Next is the robot. DAT and MOD are fast, reasonably priced, and magnetic. Magnetic media allows you to a) migrate the data (so Mrs. Jones' film from last month is written near Mrs. Jones' film from two years ago - making prefetching faster and more reliable; and b) refresh the media (replace tapes or MODs that are wearing out).

Now I know you're stuck on the media. Sony and Panasonic are beckoning you with the dream of fifty-cent DVD platters that hold Rambo 12 and you want to use them for all your chest x-rays, caths, and echocardiograms, right? OK, OK. I understand. It's coming. It's cheap. And you want it. Storage management of terabytes with good response is a tough business. When you think about AT&T only having a couple of terabytes of data from their last 30 years of telephone transactions, you begin to appreciate how big the task is of the lowly 100 bed hospital. No matter who you are, you gotta be as sophisticated as AT&T about this stuff or some lawyer may eat your lunch. So when you shop, first focus on the software (see below) and the ***reliability and speed*** of the hardware. Reliability is tough to predict on DVD jukeboxes. No one knows yet. Certainly their speed is good for movies but maybe not so good in the 'get it now' environment of morning x-ray rounds.

Now most importantly a look at the HSM software. HSM is geek-talk for hierarchical storage management. It means the software, which manages the data and knows where all of the goodies live. This is no small task when people have lots of studies separated in time and often in space (multiple institutions). Good software will not only manage different types of media, but scale from small sizes (100 gigabytes) to gigantic proportions (multiple terabytes). It will migrate data and manage refreshing the media.

Often, departments will elect to start with a 'shelf' management system. This is pretty much like the tape racks that those of us who are over 40 remember came with our original CT scanners. You scan, archive to tape, put the tape on the shelf. Still, you need some HSM software to tell you which tape, MOD or platter Mrs. Jones' study is on.

The next step usually occurs once people start using workstations for reading and or review. It takes too much time and labor to find all those studies, re-mount them on the host CT, MR, Nukes or Ultrasound device and send them to the workstation. So instead of archiving to a shelf, you start to archive to a robot or (if you're smart) to a RAID with HSM software that automatically migrates the data to the longer term robot after user-specified time intervals or after failure of users to access the study after a user specified time interval.

So before you know it, you're in a three-tiered archive environment. Then even wilder things can happen. Let's say your hospital or imaging center gets 'integrated' or bought or merged with another hospital chain or facility. All of a sudden there's one admit-discharge-transfer system, everyone wants a unified medical record and patient environment so the patients can get care in 26 medical offices and 3 hospitals all across the new empire. Now you and the hospital down the road gotta share an archive, or (worse yet) migrate your current (and no doubt different) archives on to a new mega-archive in a central basement or bomb-shelter in the middle of town. Who ya gonna call, Fred's Audio & Video Emporium?

The CD jukebox you buy for your little imaging center will probably be just fine if your doctors are tolerant and patient people and if your imaging center never gets purchased or integrated into a multi-institutional or hospital environment and if the next wave of CD/DVD readers doesn't antiquate your archive like an 8-track tape collection. I advise you to demand software which will (or you will pay through the nose to) migrate and convert your data when you hit the top of the supported archive size.

Most HSM software will not support CD or DVD yet. But very soon they will and you'll be able to avail yourself of it a year or two after you get one for your home. Some HSM software will not scale from shelf to multi-terabyte proportions. Nor will some HSM software permit a 'distributed architecture' - that's geekspeak for software that can locate and keep track of studies in multiple physical locations (e.g. your chest x-ray on the hospital jukebox, your MRI on the imaging center MOD drive, and your bunion films in the file room at the MSO facility).

Critical to this type of archive is an 'industrial grade' or enterprise scale database. Little old Microsoft SQL is pretty cool for most things, but it just won't do once you get past 100 gigabytes or so. So we're talking Informix, Oracle, and Sybase here. While you don't need to drill down to all these gory details, when it comes to a departmental archive you will in general get what you pay for.   Please remember, there's no free lunch; you can't build a skyscraper on a wood frame; the latest is the most untested solution and you can't afford failure or down-time when it comes to archives.

Conclusion

The moral to this story is as follows:

In the next few years you will have the exciting opportunity to get involved in electronic imaging and PACS in radiology. This is an exciting time, but it is different in many ways from buying CT scanners and fluoro rooms. (Indeed, buying those things is going to change as a result of PACS as we have seen above.) You are really engaging in buying components to a computer network. These components will need to inter-operate to accommodate demands which already exist, demand which will be created once you have your images available on request, and demands none of us is able to foresee at this time.

None of this endeavor is a unique or new thing. Networking IS people have been struggling with these issues for decades and have a burgeoning, exciting industry to show for it. All that is unique is that our information is in very large files, impacts on the health of a patient, and is regulated by the Food & Drug Acts.

In the near future you will work with several PACS and Teleradiology vendors, system integrators, and traditional device companies. Try to learn from them, while keeping some perspective on what you are doing and what you wish to accomplish. If you can, try to think two or three steps down the road. No one vendor is likely to have designed, developed and, as a result, be able to deliver everything you will need. It isn't the case in the computer networking environment and it won't be true in the PACS and Teleradiology environment either. Try to stick to each vendor's core competencies and demand of them interoperability and conformance to DICOM, HL-7, SQL standards and any other standards, which become pertinent along the way.

Most of all, get wired. Your patients and customers will need you to be.

Comments and Questions are encouraged to: Phillip Berman, M.D. (berman@lumisys.com)